Modeling Time Series Similarity with Siamese Recurrent Networks

Wenjie Pei W.PEI-1@TUDELFT.NL
David M.J. Tax D.M.J.TAX@TUDELFT.NL
Delft University of Technology, Mekelweg 4, 2628 CD Delft, THE NETHERLANDS

Laurens van der Maaten LVDMAATEN@FB.COM
Facebook AI Research, 770 Broadway 8th Floor, New York NY 10003, USA

arXiv:1603.04713v1 [cs.CV] 15 Mar 2016

Abstract

Traditional techniques for measuring similarities between time series are based on handcrafted similarity measures, whereas more recent learning-based approaches cannot exploit external supervision. We combine ideas from time-series modeling and metric learning, and study *siamese recurrent networks* (SRNs) that minimize a classification loss to learn a good similarity measure between time series. Specifically, our approach learns a vectorial representation for each time series in such a way that similar time series are modeled by similar representations, and dissimilar time series by dissimilar representations. Because it is a similarity prediction models, SRNs are particularly well-suited to challenging scenarios such as signature recognition, in which each person is a separate class and very few examples per class are available. We demonstrate the potential merits of SRNs in within-domain and out-of-domain classification experiments and in one-shot learning experiments on tasks such as signature, voice, and sign language recognition.

1. Introduction

Successful classification, verification, or retrieval of time series requires the definition of a good similarity measure between time series. Classical approaches to time-series analysis handcraft such similarity measures (Vintsyuk, 1968; Sakoe & Chiba, 1978), which limits their ability to incorporate information on the relative scale of features in the similarity measure. Other approaches use unsupervised learning in order to define the similarity measure (Rabiner, 1989; Jaakkola & Haussler, 1998), which has the disadvantage that it cannot exploit class label information in determining which features are most relevant for the underlying similarity structure.

In this paper, we study a novel model for time-series analysis that *learns* a similarity measure over pairs of time series in a *supervised* manner. The proposed model combines ideas from metric learning with that of learning embeddings for time series using recurrent networks. The model takes as input two time series, which are both processed by the same recurrent network to produce a representation for each of time series. The similarity between the time series is defined as a weighted inner product between the resulting representations. All parameters of the model are learned jointly by minimizing a classification loss on pairs of similar and dissimilar time series. We refer to the resulting model as *siamese recurrent network* (SRN). The structure of the SRN is illustrated in Figure 1. We evaluate the performance of two variants of the SRN in within-domain classification and out-of-domain classification experiments representing a range of different machine-learning tasks.

The model we study in this paper is of particular interest in challenging learning settings in which the number of classes is large and the number of training examples per class is limited. An example of such a setting is an online signature verification task. Here each person who provided one or more signatures is considered to be a separate class, and the number of training examples per person is extremely limited. Such a task may benefit from sharing parameters between classes by learning a global similarity measure over the set of all pairs of time series, which is what the SRN does. We perform one-shot learning experiments to illustrate the potential merits of the global similarity measure over time series learned by our models.

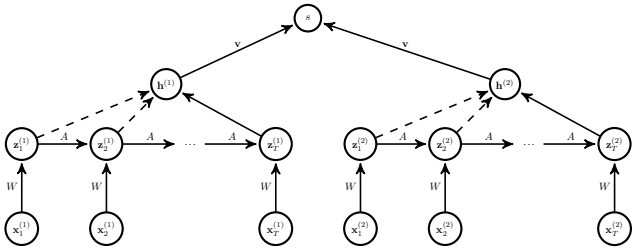

Figure 1. Graphical representation of the Siamese Recurrent Network (SRN). For the SRN-L model, the feature representations $\mathbf{h}$ are obtained by taking the hidden unit activations at the last timestep, $\mathbf{z}_T$ (solid line). For the SRN-A model, the feature representations $\mathbf{h}$ are obtained by averaging the hidden unit activations $\mathbf{z}$ over all timesteps (solid and dashed lines). The SRN outputs a scalar similarity measure s.

2. Related Work

Traditional approaches to measuring time-series similarity such as dynamic time warping (DTW; Vintsyuk (1968); Sakoe & Chiba (1978)) use handcrafted similarity measures that are not adapted to the observed data distribution. This shortcoming was addressed by the introduction of similarity measures that first fit a generative model to the data, such as Fisher, TOP, marginalized, and product-probability kernels (Jaakkola & Haussler, 1998; Tsuda et al., 2002a; Jebara et al., 2004; Tsuda et al., 2002b). In particular, Fisher kernels have seen widespread adoption in computer vision (Perronnin et al., 2010). While these methods benefit from modeling the data distribution before the computation of pairwise similarities, they are limited in that they cannot exploit available supervised class or similarity information, which may hamper their performance in classification problems. By contrast, the time-series similarity approach we study in this work is based on *supervised learning*. It combines ideas from modeling time series using recurrent networks with those from metric learning. We discuss related work on both topics separately below.

Recurrent networks learn a representation for each timestep that is influenced by both the observation at that time step and by the representation in the previous timestep (Werbos, 1988; Schmidhuber, 1989). The recurrent nature of the models equips them with a memory that is capable of preserving information over time. This has made them popular for tasks such as language (Mikolov et al., 2011; Vinyals et al., 2015), handwriting (Graves, 2013), and image generation (Theis & Bethge, 2015), and music prediction (Bengio et al., 2013). SRNs employ a pair of standard recurrent networks, the parameters of which are shared between the two networks. It differs from prior work in the loss that it minimizes: instead of minimizing a "generative" loss such as negative log-likelihood, it minimizes a loss that encourages representations to be close together for similar time series and far apart for dissimilar time series.

Metric learning techniques learn a similarity measure on data that lives in a vectorial space. While several studies have explored learning non-linear "metrics" by back-propagating pairwise losses through feedforward networks (Bromley et al., 1993; Chopra et al., 2005; Salakhutdinov & Hinton, 2007; Koch et al., 2015; Min et al., 2010; Hadsell et al., 2006; Hu et al., 2014), most prior work on metric learning focuses on learning Mahalanobis metrics; prominent examples of such studies include Goldberger et al. (2004); Weinberger & Saul (2009); Davis et al. (2007); and Xing et al. (2002). Our work is most similar to latent coincidence analysis (LCA; Der & Saul (2012)) in terms of the loss it is minimizing, but it differs substantially from LCA in that it backpropagates the loss through the recurrent network that is modeling the time series.

3. Siamese Recurrent Networks

A time-series similarity model produces a single similarity value for each input pair of time series (with potentially different lengths). Similarly to a siamese network, our time-series similarity model employs two neural networks that share their network parameters in order to extract comparable hidden-unit representations from the inputs. The resulting hidden-unit representations are compared to compute the similarity between the two time series. The parameters of the neural networks and the comparison function are learned jointly in a supervised manner to predict whether two time series are similar or not. We use recurrent networks as the basis for our siamese architecture, leading to the siamese recurrent network (SRN) depicted in Figure 1. The advantage of using recurrent networks is that they allow our model (1) to extract relevant features for the similarity computation and (2) to remember these relevant features over time when needed. The resulting features have the same size irrespective of the time series length.

Suppose we are given two time series $X^{(1)} = \left\{\mathbf{x}_1^{(1)}, \ldots, \mathbf{x}_{T_1}^{(1)}\right\}$ and $X^{(2)} = \left\{\mathbf{x}_1^{(2)}, \ldots, \mathbf{x}_{T_2}^{(2)}\right\}$ whose lengths are respectively T_1 and T_2. The hidden-unit representations $\mathbf{z}_t^{(1)}$ and $\mathbf{z}_t^{(2)}$ in the SRN model are defined as:

$$\mathbf{z}_t^{(i)} = \mathrm{g}\left(\mathbf{W}\mathbf{x}_t^{(i)} + \mathbf{A}\mathbf{z}_{t-1}^{(i)} + \mathbf{b}\right). \tag{1}$$

We use a rectified linear unit (ReLU) function $\mathrm{g}(x) = \max(0, x)$ as this activation function eliminates potential vanishing-gradient problems.

The hidden-unit representations obtained from the two sub-networks for the corresponding input time series, $\mathbf{h}^{(1)}$ and $\mathbf{h}^{(2)}$, are combined to compute the SRN's prediction for the similarity of the two time series. We consider two approaches for comparing hidden-unit representations.

In the first approach, the element-wise product between the hidden representations on the last time steps T_1 and T_2 is computed and the output is a weighted sum of the resulting products. This approach encourages the recurrent networks to remember relevant features over time, thereby making these features available for the final similarity computation.

In the second approach, all the hidden-unit representations for each of the two time series are averaged over time to construct a single feature representation for both time series, and the resulting feature representations are combined in the same way as before to compute the time-series similarity. This approach removes the burden on the recurrent networks to memorize all important features over time, but may potentially pollute the time-series features by averaging over time.

Mathematically, the two approaches compute the following latent representations $\mathbf{h}$ for each time series:

- The **SRN-L** (last timestep) model:

$$\mathbf{h}^{(i)} = \mathbf{h}\left(X^{(i)}\right) = \mathbf{z}_T^{(i)}. \tag{2}$$

  The recurrent connections in recurrent networks allow it to memorize the previous inputs in the hidden states in a recursive way. Consequently, the hidden units in the last time step should be able to store the information accumulated in the time domain for the whole time series. Therefore, we conjecture it is capable of modeling the entire time series.

- The **SRN-A** (average) model:

$$\mathbf{h}^{(i)} = \mathbf{h}\left(X^{(i)}\right) = \frac{1}{T}\sum_{t=1}^{T}\mathbf{z}_t^{(i)}. \tag{3}$$

  By averaging the hidden units $\mathbf{z}$ over time, this model treats the information of each time step equally and avoids the potential memory-vanishing problem whilst still considering the temporal information in the previous time steps when computing hidden-unit representations.

Denoting the latent representations obtained from the two recurrent networks as $\mathbf{h}^{(1)}$ and $\mathbf{h}^{(2)}$, the SRN model defines the similarity of the two time series as:

$$s\left(X^{(1)}, X^{(2)}\right) = \frac{1}{1 + e^{-\mathbf{v}^\top\left[\mathrm{diag}\left(\mathbf{h}^{(1)}\mathbf{h}^{(2)\top}\right)\right] + c}}. \tag{4}$$

Herein, the similarity between two time series is defined as a weighted inner product between the latent representations $\mathbf{h}^{(1)}$ and $\mathbf{h}^{(2)}$. Such similarity measures between hidden-units activations have previously been used as part of attention mechanisms in speech recognition (Chorowski et al., 2014), machine translation (Bahdanau et al., 2014), and handwriting generation (Graves, 2013).

3.1. Parameter Learning

Suppose we are given a training set $\mathcal{T}$ containing two sets of in total N pairs of time series, a set with pairs of similar time series $\mathcal{S}$ and a set with pairs of dissimilar time series $\mathcal{D}$. We learn all parameters $\Theta = \{\mathbf{A}, \mathbf{W}, \mathbf{v}, c, \mathbf{b}\}$ of the SRN jointly by minimizing the binary cross-entropy of predicting to which set each pair of time series belongs with respect to the parameters. This is equivalent to maximizing the conditional log-likelihood of the training data:

$$\mathcal{L}(\Theta; \mathcal{T}) = -\Bigg[\sum_{(n_1,n_2)\in\mathcal{S}} \log s\left(X^{(n_1)}, X^{(n_2)}\right) + \sum_{(n_1,n_2)\in\mathcal{D}} \log\left(1 - s\left(X^{(n_1)}, X^{(n_2)}\right)\right)\Bigg],$$

where n_1 and n_2 indicate the indices of the first and second time series in a training pair. The loss function is back-propagated through both recurrent networks (the weights of which are shared) using a variant of the backpropagation through time algorithm (Werbos, 1988) with gradient clipping between -5 and 5 (Bengio et al., 2013).

The sets $\mathcal{S}$ and $\mathcal{D}$ of similar and dissimilar time series can be constructed in various ways, for instance, by asking human annotators for similarity judgements. When class labels y_n are available for time series $X^{(n)}$, the sets can be defined as $\mathcal{S} = \{(n_1, n_2) : y_{n_1} = y_{n_2}\}$ and $\mathcal{D} = \{(n_1, n_2) : y_{n_1} \neq y_{n_2}\}$. In contrast to time-series classification models (Eddy et al., 1995; Kim & Pavlovic, 2006; van der Maaten, 2011; Quattoni et al., 2010), this allows SRNs to be used on objects from unknown classes as well. For instance, the SRN may be trained on the signatures of a collection of people, and like any classification model, it can then be used within-domain to verify new signatures of the same people. However, the SRN can also be used out-of-domain to verify the signatures from people that were not present in the training set. The SRN only needs one genuine, verified signature to compute the similarity to a new, unknown signature (one-shot learning). The underlying assumption is that the inter-person variation of the signatures is modeled well by the SRN because it was trained on signatures from many other people.

4. Experiments

We performed experiments with SRNs on three different datasets in three different learning settings: (1) within-domain similarity prediction, (2) out-of-domain similarity prediction, and (3) one-shot learning. Before presenting the setup and results of our experiments, we first introduce the three datasets below.

Table 1. Characteristics of the five datasets considered in our experiments: dimensionality of features, number of classes, number of samples, and the minimum, mean, and maximum length of the time series.

Dataset	Dimens.	Classes	Samples	Time series length		
				Min.	Mean	Max.
Arabic (digit)	13×2	10	8800	3	39	92
Arabic (voice)	13×2	88	8800	3	39	92
MCYT (without forgery)	5×3	100	2500	34	349	1161
MCYT (with forgery)	5×3	100	5000	34	438	2687
Sign	77×2	19	40	760	112	198

4.1. Datasets

We performed experiments on three different datasets.

The *Arabic Spoken Digit dataset* (Hammami & Sellami, 2009) comprises $8,800$ utterances of digits produced by 88 different speakers. Each speaker uttered each digit ten times. The data is represented as a time series of 13-dimensional MFCCs that were sampled at $11,025$Hz and 16 bits using a Hamming window. We use two different versions of the spoken digit dataset: (1) a *digit* version in which the uttered digit is the class label and (2) a *voice* version in which the speaker of a digit is the class label.

The *MCYT signature dataset* (Ortega-Garcia et al., 2003) contains online signature data collected from 100 subjects. For each subject, the data comprises 25 authentic signatures and 25 skilled forgeries. The signatures are represented as time series of five features: the x-coordinate, y-coordinate, pressure, azimuth, and elevation of the pen. We consider two different versions of the dataset, namely, a version *without forged* data and a version *with forged* data.

The *American sign language dataset* (Aran et al., 2006) contains eight manual signs that represent different words and eleven non-manual signs such as head or shoulder motions. The data thus comprises nineteen classes. Each sign was produced five times by eight different subjects, leading to a total of 760 samples. The time series are represented using a hand-crafted feature representation that contains a total of 77 hand motion, hand position, hand shape, and head motion features (Aran et al., 2006).

Following common practice in time series analysis, we preprocessed all three datasets by applying a sliding window (with stride 1) to the time series, concatenating the features in the frames under the window into a single frame. This enriches the feature representation, making it easier for the models to capture feature gradients. For the Arabic, MCYT, and Sign datasets, we used a window size of 2, 3, and 2, respectively. In Table 1, the main characteristics of all five datasets are summarized.

4.2. Experimental setup

In our experiments, the model parameters of the SRNs were initialized by sampling them from a uniform distribution within an interval $[-0.1, 0.1]$. Training of the model is performed using a RMSprop (Tieleman & Hinton, 2012) stochastic gradient descent procedure using mini-batches of 50 pairs of time series. To prevent the gradients from exploding, we clip all gradients (Bengio et al., 2013) to lie in the interval $[-5, 5]$. We decay the learning rate during training by multiplying it by 0.4 every time the AUC on the validation set stops increasing. We applied dropout on the hidden-unit activations of our model: the dropout rate was tuned to maximize the AUC on a small held-out validation set. Code reproducing the results of our experiments is available on `http://www.anonymized.com`.

In all experiments except for those on the MCYT (with forgery) dataset, we defined the sets of similar and dissimilar time series as suggested in Section 3, that is, we define similar time series to be those with the same class label and dissimilar time series to be those with different class labels: $\mathcal{S} = \{(n_1, n_2) : y_{n_1} = y_{n_2}\}$ and $\mathcal{D} = \{(n_1, n_2) : y_{n_1} \neq y_{n_2}\}$. Herein, y_n represents the class label of the time series as described in section 4.1. On the MCYT (with forgery) dataset, we define the positive pairs in the same way but we define the set of negative pairs $\mathcal{D}$ slightly differently: the negative pairs are pairs of a genuine signature and a forged version of the same signature. These negative pairs are more difficult to distinguish, as a result of which training on them will likely lead to better models.

We compare the performance of our SRNs with that of three variants of our model, and with three baseline models. The three variants of our model we consider are: (1) a feedforward variant of SRN-A, called SN-A, that removes all recurrent connections from the model, *i.e.*, in which $\mathbf{A} = \mathbf{0}$ but which still averages the hidden representation over time; (2) a feedforward variant of SRN-L, called SN-L, that removes all recurrent connections from the model and uses the hidden representation of the last time step; and (3) a naive logistic model that removes all hidden units from the model and that predicts similarities by averaging all features over time and computing a weighted sum of the element-wise product of the resulting feature representations. These three variants of SRNs allow us to investigate the effect of the recurrent connections and non-linearities on the prediction performance of our models.

The three time-series similarity models we use as baseline models are: (1) dynamic time warping (Vintsyuk, 1968); (2) Fisher kernels (Jaakkola & Haussler, 1998); and (3) Fisher vectors (Perronnin et al., 2010). Details of these three baseline models are given below.

Dynamic time warping (DTW; Vintsyuk (1968)) mea-

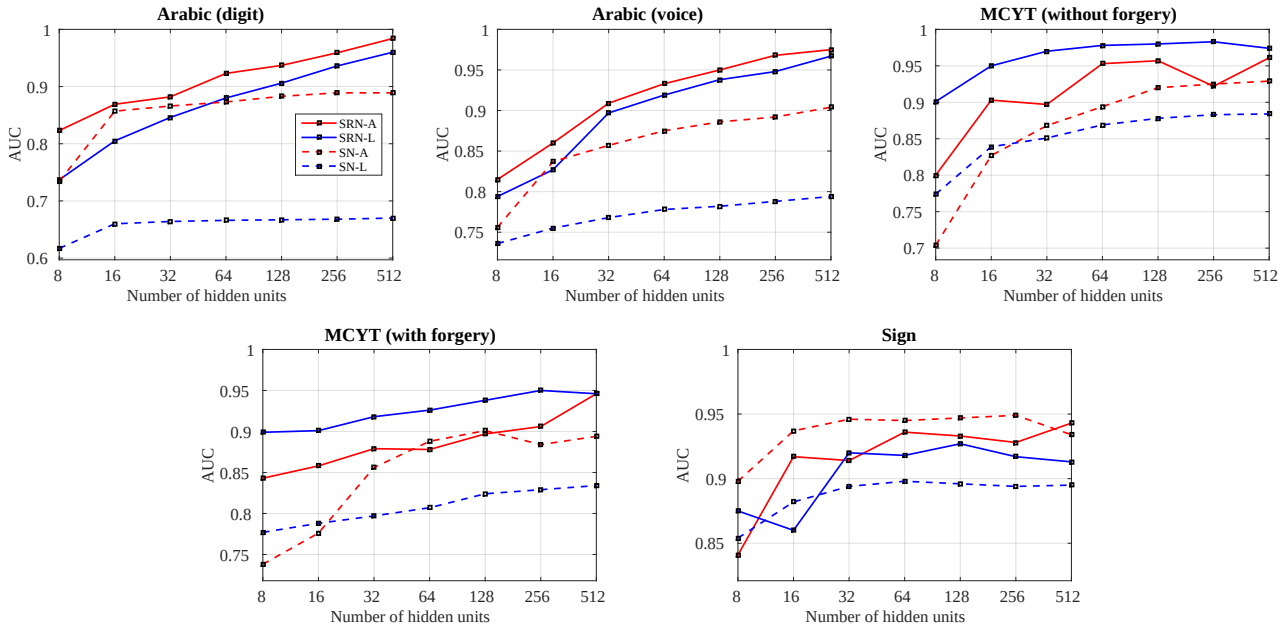

Figure 2. Area under the receiving-operator curve curve (AUC) of our two variants of Siamese Recurrent Networks (SRN-A and SRN-L) on five datasets as a function of the number of hidden units (higher is better). For reference, the performance of SRNs without recurrent connections (SNs) is also shown. All results were obtained by averaging over five repetitions. The standard deviation of the results is typically smaller than 0.01.

sures time series similarities by aligning both time series and summing the pairwise distances between all corresponding frames, minimized over the set of all possible alignments between the two time series. An alignment is a set of (potentially many-to-many) correspondences between frames, with the restriction that correspondences cannot be crossing each other in time. DTW similarities can be computed efficiently using a dynamic-programming algorithm. Despite its simplicity, DTW has been quite successful, in particular, on problems in which the time series are already relatively well aligned and the time series show some clear salient features in time. We leave comparisons with approaches that combine dynamic time warping and metric learning (Garreau et al., 2014) to future work.

Fisher kernels measure the similarity between two time series by the inner product of the log-likelihood gradients that are induced by the time series with respect to the parameters of a generative model (Jaakkola & Haussler, 1998). Our generative model of choice for time series is the hidden Markov model (HMM). Mathematically, we denote the gradient of the log-likelihood $L(X^{(n)})$ of a time series $X^{(n)}$ with respect to the parameters of the HMM as $\mathbf{g}_n = \left[\forall\theta \in \Theta : \frac{\partial L(X^{(n)})}{\partial\theta}\right]$. We define the Fisher kernel similarity κ between two time series as an inner product between their corresponding gradients:

$$\kappa\left(X^{(i)}, X^{(j)}\right) = \mathbf{g}_i^\top \mathbf{U}^{-1} \mathbf{g}_j. \tag{5}$$

Herein, the matrix $\mathbf{U}$ is the Fisher information metric, which is replaced with identity matrix $\mathbf{I}$ in our experiments. The number of hidden states of our HMMs is tuned by maximizing the AUC on a small, held-out validation set.

Fisher vectors compute the same gradients $\mathbf{g}_n$ as before, but instead of computing their inner products, we concatenate the gradients $\mathbf{g}_i$ and $\mathbf{g}_j$ to obtain a feature representation of the time-series pair $\left(X^{(i)}, X^{(j)}\right)$. Such Fisher vector representation are commonly used in computer vision (Perronnin et al., 2010). Because the concatenated Fisher vectors cannot directly measure time-series similarity, we perform 1-nearest classification on the collection of similar and dissimilar pairs to predict whether a pair of time series is similar. (In other words, the time series similarity is the negative Euclidean distance between the example and its nearest pair of similar time series in the concatenated Fisher vector space.)

4.3. Results

Below, we separately present the results for the three learning settings we considered: (1) within-domain similarity

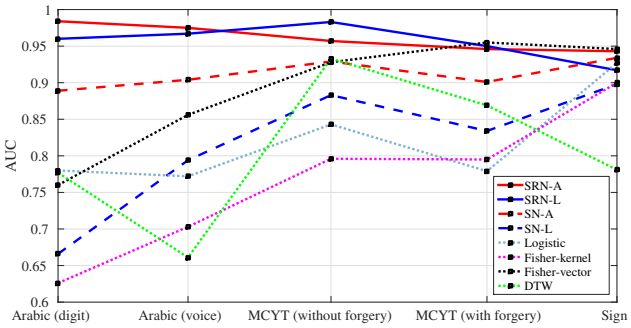

Figure 3. Area under the receiving-operator curve curve (AUC) on five different datasets using eight different time-series similarity learning models in a within-domain similarity prediction setting (higher is better). See text for details.

prediction, (2) out-of-domain similarity prediction, and (3) one-shot learning. We also present t-SNE visualizations of the learned time-series representations.

4.3.1. Within-Domain Similarity Prediction

We first evaluate the within-domain similarity prediction performance of the SRN: we randomly split the time series into a training and a test set, and we measure the ability of the models to accurately predict whether pairs of time series in the test set are similar or not in terms of the area under the receiving-operator curve (AUC). We opt for the AUC as a performance measure because it naturally deals with the potential imbalance in the sizes of $\mathcal{S}$ and $\mathcal{D}$. We refer to this experiment as *within-domain* because all classes in the test data were also observed during training.

Figure 2 presents the within-domain similarity prediction performance of SRNs as a function of the number of hidden units in the model on five different datasets. We present results for both the variant that averages all hidden-unit activations over time (SRN-A) and the variant that uses only the hidden unit activations at the last timestep (SRN-L). The reported results were averaged over five repetitions, randomly initializing the parameter of the models in each repetition. The figure also reports the performance of models without recurrent connections, called a Siamese network (SN, where SN-A is a Siamese network with averaged hidden activations and SN-L is a network that uses the last time step activations). From the results presented in Figure 2, we make three main observations.

First, the results show that the performance of SRNs tends to increase with the number of hidden units, in particular, on challenging datasets such as the Arabic speech datasets. This shows that SRNs effectively use the additional capacity that is provided by additional hidden units to learn more informative features for the time-similarity measurements. In our experiments, we did not observe much overfitting, although overfitting is likely to occur when the number of hidden units is increased much further.

Second, we observe that there is no clear winner between averaging hidden unit activations over time (SRN-A) and using the activations at the last timestep (SRN-L). This suggests that the recurrent networks in the SRN-L models are at least partly successful in remembering relevant features over time.

Third, we observe that the recurrent connections in the SRN models are, indeed, helpful: the SRN models outperform their counterparts without recurrent connections (SNs) in nearly all experiments[1]. This result underlines the hypothesis that recurrent connections can preserve features relevant for time-series similarity computations over time. Somewhat surprisingly, the performance of the SN-L models is not as bad as one may expect. It should be noted that the windowing of features makes the feature representation of the last timestep richer, which is sufficient to obtain acceptable performances on some of the datasets.

Comparison with baseline models. Next, we compare the performance between of SRNs with the naive logistic model and three other baseline time-series similarity learning models: (1) dynamic time warping, (2) Fisher kernels, and (3) Fisher vectors (see section 4.2 for details). We used the same experimental setup as in the previous experiment, but we tuned the main hyperparameters of the models (the number of hidden units in SRNs and SNs; the number of HMM hidden states for Fisher kernels and Fisher vectors) on a small held-out validation set. Figure 3 presents the results of these experiments.

The results of these experiments show that, indeed, the SRN can be a very competitive time-series similarity model, even when trained on relatively small datasets. In particular, SRNs substantially outperform the baselines models on the Arabic (digit), Arabic (voice), and MCYT (without forgery) datasets. On most datasets, the Fisher vectors are the best baseline model (they perform substantially better than standard Fisher kernels), which is line with results in the literature (Perronnin et al., 2010). The naive logistic model performs substantially worse than the SRN models, which suggests that hidden units are essential in solving difficult similarity assessment problems.

Dynamic time warping (DTW) performs reasonably well on relatively simple datasets such as the Sign dataset, but its performance deteriorates on more challenging datasets in which the similarity labels are not aligned with the main

[1]It should be noted that because we preprocess the time-series data by windowing features, the SN is actually a convolutional network that is very similar to the time-delay neural networks of Bromley et al. (1993).

Table 2. Area under the receiving-operator curve curve (AUC) of eight time-series similarity models on five datasets in an out-of-domain similarity prediction setting (higher is better). The standard deviation of the five repetitions we performed is typically smaller than 0.01. The best performance per dataset is boldfaced. See text for details.

Dataset	Training classes	Test classes	Model							
			SRN-A	SRN-L	SN-A	SN-L	Logist.	DTW	Fisher K.	Fisher V.
Arabic (digit)	1-7	8-10	0.681	0.714	**0.768**	0.539	0.761	0.725	0.600	0.561
Arabic (voice)	1-60	61-88	**0.849**	0.788	0.802	0.684	0.730	0.640	0.698	0.630
MCYT (without forgery)	1-70	71-100	0.914	0.920	0.816	0.760	0.824	**0.952**	0.752	0.844
MCYT (with forgery)	1-70	71-100	0.888	0.876	0.828	0.668	0.782	**0.894**	0.805	0.813
Sign	1-14	15-19	**0.862**	0.670	0.748	0.565	0.836	0.729	0.770	0.566

sources of variation in the data, such as the Arabic (voice) dataset: the main sources of variation in this dataset are likely due to the differences in the digits being uttered, whereas the similarity labels we are interested in concern the speaker of the digit and not the digit itself. DTW (as well as Fisher vectors and kernels) cannot exploit this information, which explains its inferior performance on the Arabic (voice) dataset.

4.3.2. Out-of-Domain Similarity Prediction

In the next set of experiments, we measure the performance of SRNs on *out-of-domain* similarity prediction: we use the same experimental setup as before, however, we split the training and test data in such a way that the set of class labels appearing in the training set and the set of class labels appearing in the test set are disjoint. This is a more challenging learning setting, as it relies on the time-series similarity model exploiting structure that is shared between classes in order to produce good results. We obtain the test data by selecting 3 out of 10 classes on the Arabic (digit) dataset, 28 out of 88 classes on the Arabic (voice) dataset, 30 out of 100 classes on the MCYT datasets, and 5 out of 19 classes on the Sign dataset. As before, we measure the performance of our models in terms of AUC, and we tune the hyperparameters of the models on a validation set. The results of these experiments are presented in Table 2.

From the results presented in the table, we observe that the strong performance of SRNs on difficult datasets such as the Arabic (voice) datasets generalizes to the out-of-domain prediction setting. This suggests that, indeed, the SRN models are able to learn some structure in the data that is shared between classes. On the (much smaller) MCYT datasets, however, dynamic time warping outperforms SRNs. Most likely, this result is caused by the SRNs (which have high capacity) overfitting on the classes that are observed during training.

Table 3. Classification accuracy of one-shot learning models of an 1-nearest neighbor classifier using three different similarity measures on four different datasets (higher is better). The best performance per dataset is boldfaced. See text for details.

Dataset	SRN-A	SRN-L	DTW
Arabic (digit)	0.618	0.613	**0.801**
Arabic (voice)	**0.273**	0.228	0.151
MCYT (without forgery)	0.418	0.548	**0.913**
Sign	**0.599**	0.381	0.531

4.3.3. One-Shot Learning

To further explore the potential of SRNs in out-of-domain settings, we performed an experiment in which we measured the performance of SRNs in one-shot learning. We adopt the same dataset splits as in 4.3.2 to obtain train and test data. On the training data, we train the SRNs to learn a similarity measure for time series. This similarity measure is used to train and evaluate a nearest-neighbor classifier on the test set. We use only a single time series per class from the test set to train the nearest-neighbor classifier, and use the remaining time series in the test set for evaluation. We measure the classification accuracy using leave-one-per-class-out validation.

The results are presented in Table 3. For datasets that have clear salient features, like the MCYT, and to a lesser degree the Sign dataset, DTW performs well. For more complex data, the SRN performs well provided that sufficient training data is available. For the Arabic (digit) dataset, the seven classes used in training are insufficient for the SRN, and the SRN overfits on those classes. On the Arabic (voice) dataset 60 classes are available, which allows the SRN to fully exploit its potential.

4.3.4. Visualizing the Representations

The one-shot learning experiment presented above exploits an interesting property of the SRN model, namely, that it learns a single embedding for a time series. An advantage of this is that the resulting time-series embeddings can be used in a wide variety of other learning algorithms that operate on vectorial data, such as alternative classification techniques, clustering models, *etc.* To obtain more insights into what the SRN models have learned, we apply t-SNE (van der Maaten & Hinton, 2008) on embeddings obtained by a SRN-L on the MCYT (without forgery) test set. Figure 4 shows a map of the 2500 signatures in the test set; the signatures were drawn by integrating the pen movements over time. The color indicates the pen pressure. We refer the reader to the supplemental material for a full-resolution version of this map. The t-SNE visualization shows that, indeed, the SRN-L is capable of grouping similar signatures together very well.

In Figure 5, we show a t-SNE map of the Arabic (voice) test set constructed on SRN-L embeddings. For comparison, we also show a t-SNE map of the same data, based on pairwise distances computed with DTW. The two maps clearly show the potential advantage of SRN: it has used the supervised similarity information to group all the utterances corresponding to a single class together, something that DTW is unable to do due to its unsupervised nature.

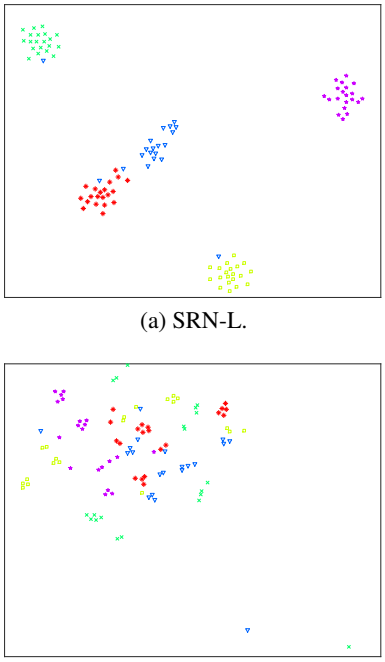

(a) SRN-L.

(b) DTW.

Figure 5. t-SNE maps of the Arabic (voice) test data from five randomly selected classes, constructed based on (a) siamese recurrent network (last timestep) embeddings of the time series and (b) pairwise similarities computed using dynamic time warping.

5. Conclusions

We have investigated models for learning similarities between time series based on supervised information. Our study shows that a combination of ideas from metric learning and deep time-series models has the potential to improve the performance of models for time-series classification, retrieval, and visualization. The proposed siamese recurrent networks (SRNs) are particularly effective compared to alternative techniques in settings in which the similarity function that needs to be learned is complicated, or when the number of labeled time series for some of the classes of interest is limited. When a reasonably large collection of examples of similar and dissimilar time series is available to train the models, the siamese recurrent networks can produce representations that are suitable for challenging problems such as one-shot learning or extreme classification of time series. This result is in line with earlier results for siamese convolutional networks by, for instance, Kamper et al. (2016).

This study is an initial investigation into learning similarities between time series, and we foresee several directions for future work. In particular, we intend to explore variants of our model architecture: (1) that employ a bilinear model to measure the similarity of the RNN representations; (2) that employ long-term short-term units (Hochreiter & Schmidhuber, 1997) or gated recurrent units (Cho et al., 2014) instead of the simple rectified linear units we are currently using; (3) that employ multiple layers of recurrent units; and (4) that have a tree structure or generic (planar) graph structure instead of the current sequential structure. The latter extension would make our models applicable to problems such as molecule classification (Riesen & Bunke, 2008). We also plan to explore improvements to our learning algorithm. In particular, our current implementation selects negative pairs of time series in a somewhat arbitrary manner: in all our experiments, we select negative examples uniformly at random for the set of all candidate negative pairs. We plan to investigate approaches that perform a kind of "hard negative mining" during learning, akin to some modern metric learning (Weinberger & Saul, 2009) and multi-modal learning (Weston et al., 2011) approaches. We also plan to study applications of SRNs in, for instance, learning word-discriminative acoustic features (Synnaeve et al., 2014).

Acknowledgments

This work was supported by AAL SALIG++.

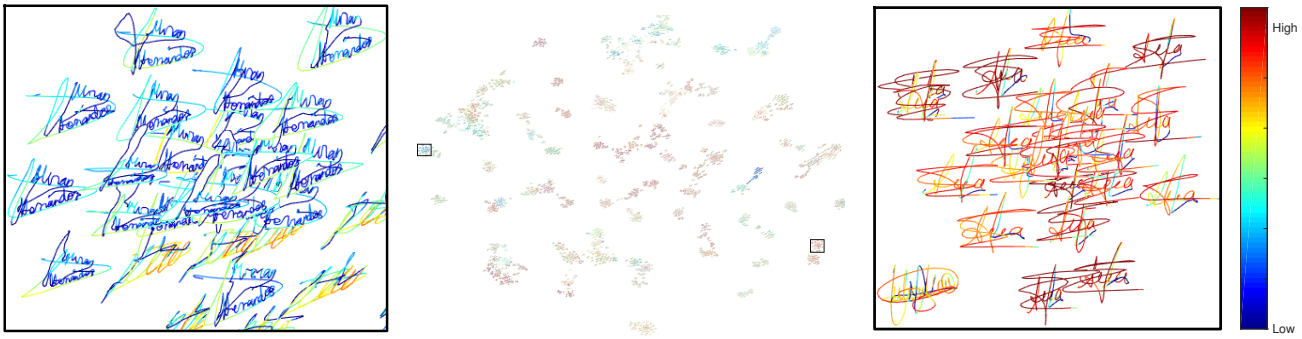

Figure 4. t-SNE map of the $2,500$ signatures in the MCYT test set (100 subjects) data based on embeddings computed by an SRN-L. The signatures were drawn by integrating the pen movements over time, and colors indicate the pen pressure (red indicates high pressure and blue indicates low pressure). A full-resolution version of this map is presented in the supplemental material.

References

Aran, O., Ari, I., Benoit, A., Campr, P., Carrillo, A. Huerta, Fanard, F.-X., Akarun, L., Caplier, A., Rombaut, M., and Sankur, B. Sign language tutoring tool. In *Proceedings of eNTERFACE 2006, The Summer Workshop on Multimodal Interfaces*, pp. 23–33, 2006.

Bahdanau, D., Cho, K., and Bengio, Y. Neural machine translation by jointly learning to align and translate. In *arXiv:1409.0473*, 2014.

Bengio, Y., Boulanger-Lewandowski, N., and Pascanu, R. Advances in optimizing recurrent networks. In *Proceedings of the IEEE Conference on Acoustics, Speech and Signal Processing*, 2013.

Bromley, J., Guyon, I., LeCun, Y., Sackinger, E., and Shah, R. Signature verification using a siamese time delay neural network. In *Advances in Neural Information Processing Systems*, volume 6, 1993.

Cho, K., van Merrienboer, B., Bahdanau, D., and Bengio, Y. On the properties of neural machine translation: Encoder-decoder approaches. In *Proceedings of SSST-8, Eighth Workshop on Syntax, Semantics and Structure in Statistical Translation*, pp. 103–111, 2014.

Chopra, S., Hadsell, R., and LeCun, Y. Learning a similarity measure discriminatively with applications to face verification. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2005.

Chorowski, J., Bahdanau, D., Cho, K., and Bengio, Y. End-to-end continuous speech recognition using attention-based recurrent NN: First results. In *Deep Learning and Representation Learning Workshop: NIPS*, 2014.

Davis, J.V., Kulis, B., Jain, P., Sra, S., and Dhillon, I. S. Information-theoretic metric learning. In *Proceedings of the International Conference on Machine Learning*, pp. 209–216, 2007.

Der, M.F. and Saul, L.K. Latent coincidence analysis: a hidden variable model for distance metric learning. In *Advances in Neural Information Processing Systems 25*, pp. 3239–3247, 2012.

Eddy, S.R., Mitchison, G., and Durbin, R. Maximum discrimination hidden Markov models of sequence consensus. *Journal of Computational Biology*, 2(1):9–24, 1995.

Garreau, D., Lajugie, R., Arlot, S., and Bach, F. Metric learning for temporal sequence alignment. In *Advances in Neural Information Processing Systems*, 2014.

Goldberger, J., Roweis, S., Hinton, G.E., and Salakhutdinov, R.R. Neighbourhood component analysis. In *Neural Information Processing Systems*, 2004.

Graves, A. Generating sequences with recurrent neural networks. In *arXiv:1308.0850*, 2013.

Hadsell, R., Chopra, S., and LeCun, Y. Dimensionality reduction by learning an invariant mapping. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2006.

Hammami, N. and Sellami, M. Tree distribution classifier for automatic spoken Arabic digit recognition. In *Proceedings of the IEEE ICITST09 Conference*, pp. 1–4, 2009.

Hochreiter, S. and Schmidhuber, J. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997.

Hu, J., Lu, J., and Tan, Y.-P. Discriminative deep metric learning for face verification in the wild. In *IEEE International Conference on Computer Vision and Pattern Recognition*, 2014.

Jaakkola, T. and Haussler, D. Exploiting generative models in discriminative classifiers. In *In Advances in Neural Information Processing Systems 11*, pp. 487–493, 1998.

Jebara, T., Kondor, R., and Howard, A. Probability product kernels. *Journal of Machine Learning Research*, 5(Dec): 819–844, 2004.

Kamper, H., Wang, W., and Livescu, K. Deep convolutional acoustic word embeddings using word-pair side information. In *Proceedings of the IEEE Conference on Acoustics, Speech and Signal Processing*, 2016.

Kim, M. and Pavlovic, V. Discriminative learning of mixture of Bayesian network classifiers for sequence classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 268–275, 2006.

Koch, G., Zemel, R., and Salakhutdinov, R. Siamese neural networks for one-shot image recognition. In *ICML 2015 Deep Learning Workshop*, 2015.

Langley, P. Crafting papers on machine learning. In Langley, Pat (ed.), *Proceedings of the International Conference on Machine Learning*, pp. 1207–1216, 2000.

Mikolov, T., Kombrink, S., Burget, L., Cernock, J., and Khudanpur, S. Extensions of recurrent neural network language model. In *ICASSP*, pp. 5528–5531. IEEE, 2011.

Min, R., van der Maaten, L.J.P., Yuan, Z., Bonner, A., and Zhang, Z. Deep supervised t-distributed embedding. In *Proceedings of the International Conference on Machine Learning*, pp. 791–798, 2010.

Ortega-Garcia, J., Fierrez-Aguilar, J., Simon, D., Gonzalez, J., Faundez-Zanuy, M., Espinosa, V., Satue, A., Hernaez, I., Igarza, J.-J., Vivaracho, C., Escudero, D., and Moro, Q.-I. MCYT baseline corpus: A bimodal biometric database. *IEE Proceedings on Vision, Image and Signal Processing*, 150(6), 2003.

Perronnin, F., Sánchez, J., and Mensink, T. Improving the Fisher kernel for large-scale image classification. In *Proceedings of the European Conference on Computer Vision*, pp. 143–156, 2010.

Quattoni, A., Wang, S., Morency, L.-P., Collins, M., and Darrell, T. Hidden conditional random fields. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 29(10), 2010.

Rabiner, L. A tutorial on hidden markov models and selected applications in speech recognition. *Proceedings of the IEEE*, 77(2):257–286, 1989.

Riesen, K. and Bunke, H. Iam graph database repository for graph based pattern recognition and machine learning. In *Lecture Notes in Computer Science*, volume 5342, pp. 287–297, 2008.

Sakoe, H. and Chiba, S. Dynamic programming algorithm optimization for spoken word recognition. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 26: 43–49, 1978.

Salakhutdinov, R.R. and Hinton, G.E. Learning a nonlinear embedding by preserving class neighbourhood structure. In *Proceedings of the International Conference on Artificial Intelligence and Statistics*, volume 11, 2007.

Schmidhuber, J. A local learning algorithm for dynamic feedforward and recurrent networks. *Connection Science*, 1(4):403–412, 1989.

Synnaeve, G., Schatz, T., and Dupoux, E. Phonetics embedding learning with side information. In *IEEE Spoken Language Technology Workshop*, pp. 106–111, 2014.

Theis, L. and Bethge, M. Generative image modeling using spatial lstms. In *Advances in Neural Information Processing Systems 28*, Jun 2015.

Tieleman, T. and Hinton, G.E. Lecture 6.5 - rmsprop, coursera: Neural networks for machine learning, 2012.

Tsuda, K., Kawanabe, M., Rätsch, G., Sonnenburg, S., and Müller, KR. A new discriminative kernel from probabilistic models. *Neural Computation*, 14(10):2397–2414, 2002a.

Tsuda, K., Kin, T., and Asai, K. Marginalized kernels for biological sequences. *Bioinformatics*, 18(Suppl 1):268–275, 2002b.

van der Maaten, L.J.P. Learning discriminative Fisher kernels. In *Proceedings of the International Conference on Machine Learning*, pp. 217–224, 2011.

van der Maaten, L.J.P. and Hinton, G.E. Visualizing high-dimensional data using t-SNE. *Journal of Machine Learning Research*, 9(Nov):2579–2605, 2008.

Vintsyuk, T.K. Speech discrimination by dynamic programming. *Kibernetika*, 4:81–88, 1968.

Vinyals, O., Kaiser, L., Koo, T., Petrov, S., Sutskever, I., and Hinton, G.E. Grammar as a foreign language. In *Advances in Neural Information Processing Systems 28*, pp. 2755–2763, 2015.

Weinberger, K.Q. and Saul, L.K. Distance metric learning for large margin nearest neighbor classification. *The Journal of Machine Learning Research*, 10:207–244, 2009.

Werbos, P.J. Generalization of backpropagation with application to a recurrent gas market model. *Neural Networks*, 1, 1988.

Weston, J., Bengio, S., and Usunier, N. Wsabie: Scaling up to large vocabulary image annotation. In *Proceedings of the International Joint Conference on Artificial Intelligence*, 2011.

Xing, E.P., Ng, A.Y., Jordan, M.I., and Russell, S. Distance metric learning, with application to clustering with side-information. In *Advances in Neural Information Processing Systems 16*, pp. 521–528, 2002.